

Device-independent quantum secret sharing using Mermin-type contextuality

Stefano Gogioso*

Department of Computer Science, Oxford University

(Dated: December 31, 2015)

We present a new quantum secret sharing protocol based on recent advances in Mermin-type contextuality scenarios, which is provably secure against postquantum nonsignaling attackers. It is a fundamental assumption of secret sharing protocols that not all players are trusted parties, and that some may collude amongst themselves and with eavesdroppers to break confidentiality. To this extent, quantum secret sharing introduces a new layer of security, enabling eavesdropping detection via entangled states and noncommuting observables. A more thorough security analysis, however, becomes crucial if the protocol relies on untrusted devices for its implementation: for example, it cannot be excluded that some players may collude with the device supplier. In this paper, we put recent developments in Mermin-type contextuality to work in a new quantum secret sharing protocol. The maximal contextuality (aka maximal non-locality, or zero local fraction) demonstrated by the measurement scenarios results in strong device-independent security against nonsignaling attackers — be they classical, quantum or postquantum — which can be operationally established by means of minimal statistical analysis.

PACS numbers: 03.65.Ta, 03.65.Ud, 03.67.Dd, 03.67.Hk

In contrast to other information security protocols, classical secret sharing comes with the intrinsic assumption that some participants cannot, to some extent, be trusted. A *dealer* is interested in sharing some *secret* with a number of *players*, with the caveat that the secret be revealed to the players only when all players agree to cooperate. Integrity and availability of communications is guaranteed by the existence of authenticated classical channels between dealer and players, and the protocol is only concerned with confidentiality, defined as the impossibility of recovering the secret unless the desired number of players cooperates. Quantum secret sharing introduces a new layer of security to the protocol, employing entangled states and noncommuting observables to detect eavesdropping.

The HBB quantum secret sharing (QSS) protocol [1–3] is based on the same measurement contexts of Mermin’s contextuality argument [4]: Alice and the $N - 1$ players share N qubits in the GHZ state (with respect to the computational basis associated with the Z observable), and randomly choose to measure their qubit in either of the mutually unbiased X or Y observables. Recent work [5, 6] in the field of categorical quantum mechanics has provided a complete characterization of Mermin-type contextuality (henceforth simply *Mermin contextuality*), the validity of which extends beyond quantum mechanics, to postquantum theories and more general operational theories. All contextuality scenarios have been shown to be quantum realizable [7].

In this Letter, we introduce a new QSS protocol, based on the classification of Mermin contextuality presented above. In the original HBB protocol, confidentiality is an immediate consequence of strong complementarity of the X and Z observables [3]; eavesdropping detection, on the other hand, is derived from mutual unbiased of the X and Y observables. We generalize this analysis to the new

protocol. Furthermore, we use contextuality to provide a device-independent quantum secret sharing (DIQSS) protocol which is provably secure against postquantum nonsignaling attackers, by now an established security model for quantum cryptography [8–11].

This work uses string diagrams for dagger symmetric monoidal categories [12–17]. Our formulation of the algorithm will take place in the category of finite-dimensional Hilbert spaces and completely positive maps [18, 19], but the formalism is general enough to allow for straightforward treatment of postquantum attack scenarios. Diagrams are written bottom to top; we denote quantum (or postquantum) channels by solid wires, and classical channels by dashed wires.

Finally, by \mathbb{G} we will indicate some fixed finite abelian group with D elements, and by one D -it we will denote the amount of information carried by a uniformly distributed \mathbb{G} -valued random variable (i.e. $\log_2 D$ bits).

Mermin measurement scenarios.—Consider the unnormalized N qubit GHZ state

$$|z^0 \dots z^0\rangle + |z^1 \dots z^1\rangle, \quad (1)$$

where $|z^0\rangle, |z^1\rangle$ are the eigenstates of the single qubit Pauli Z observable. If $|0\rangle, |1\rangle$ are the eigenstates of the single qubit Pauli X observable, then the GHZ state in (1) can (up to normalization) be rewritten as

$$\sum_{x_1 \oplus \dots \oplus x_N = 0} |x_1 \dots x_N\rangle, \quad (2)$$

where \oplus is addition in the group \mathbb{Z}_2 . To obtain more Mermin scenarios, we will generalize this argument from \mathbb{Z}_2 to arbitrary finite abelian groups.

Fix a finite abelian group $(\mathbb{G}, \oplus, 0_{\mathbb{G}})$ of order D , and consider a D -dimensional quantum system \mathcal{H} . Let $|g\rangle_{g \in \mathbb{G}}$

be an orthonormal basis for \mathcal{H} , the *generalized X observable*. We define the following *generalized GHZ state*:

$$|\text{GHZ}_{0_{\mathbb{G}}}^N\rangle := \bigcirc \begin{array}{c} \cdots \\ \vdots \end{array} = \sum_{g_1 \oplus \dots \oplus g_N = 0_{\mathbb{G}}} |g_1 \dots g_N\rangle. \quad (3)$$

The traditional GHZ state is recovered, in the form (2), as the special case $\mathcal{H} \cong \mathbb{C}^2$ and $\mathbb{G} = \mathbb{Z}_2$. Define the *generalized Z observable* by the following orthogonal basis:

$$|\chi\rangle := \sum_{g \in \mathbb{G}} \chi(g) |g\rangle \text{ for all } \chi \in \mathbb{G}^{\wedge}, \quad (4)$$

where $(\mathbb{G}^{\wedge}, \cdot, \mathbb{1})$ is the abelian group of multiplicative characters of \mathbb{G} (the group homomorphisms $\chi : \mathbb{G} \rightarrow S^1$). The qubit Pauli Z observable is obtained by taking $|z^0\rangle$ to be the trivial character $|\mathbb{1}\rangle = |0\rangle + |1\rangle$ of \mathbb{Z}_2 , and $|z^1\rangle$ to be the alternating character $|0\rangle - |1\rangle$ of \mathbb{Z}_2 . The more familiar form (1) of the traditional GHZ state is recovered in the generalized Z observable for $\mathbb{G}^{\wedge} = (\mathbb{Z}_2)^{\wedge} \cong \mathbb{Z}_2$:

$$|\text{GHZ}_{0_{\mathbb{G}}}^N\rangle = \sum_{\chi \in \mathbb{G}^{\wedge}} |\chi \dots \chi\rangle \quad (5)$$

The key operational property of the qubit X and Z observables that enables Mermin's contextuality proof and their use in the HBB QSS protocol is *strong complementarity*, as shown by [6] and [3] respectively. The generalized X and Z observables above are chosen to have exactly that property, and in fact exhibit the most general form strongly complementary observables can take in finite dimensions [20].

We now proceed to define Mermin measurements using the generalized GHZ states above. The full details are covered in [5]. Consider phase gates $P_{\alpha} : \mathcal{H} \rightarrow \mathcal{H}$ for the generalized Z observable:

$$P_{\alpha} = \sum_{\chi \in \mathbb{G}^{\wedge}} e^{i\alpha(\chi)} |\chi\rangle\langle\chi|, \quad (6)$$

where $\alpha \in (\mathbb{R}/2\pi\mathbb{Z})^{\mathbb{G}^{\wedge}}$. Without loss of generality, we will set $\alpha(\mathbb{1}) = 0$. Phase gates form a group $(\mathbb{P}, \oplus, P_0)$ if we define $P_{\alpha} \oplus P_{\beta} := P_{\alpha+\beta}$, and \mathbb{G} can be identified with a subgroup $(\mathbb{K}, \oplus, P_0)$ via the following injective group homomorphism:

$$g \mapsto P_g := \sum_{\chi \in \mathbb{G}^{\wedge}} \chi(g) |\chi\rangle\langle\chi|.$$

By a *Mermin measurement* $(\alpha_i)_{i=1}^N$ on a generalized GHZ state $|\text{GHZ}_{0_{\mathbb{G}}}^N\rangle$, we will mean an application of phase gates P_{α_i} to subsystems $i = 1, \dots, N$ of $|\text{GHZ}_{0_{\mathbb{G}}}^N\rangle$, followed by measurement of each subsystem in the generalized X observable. This is exemplified in Fig. (1). We further require that $\bigoplus_{i=1}^N P_{\alpha_i} = P_g \in \mathbb{K}$ for some $g \in \mathbb{G}$. The




FIG. 1. Mermin measurement of the generalized GHZ state. Generalized GHZ state denoted in white; measurements in the generalized X observable denoted in gray; phase gates denoted by boxes.

classical measurement outcomes are valued in G , and are defined by the following probability distribution:

$$\text{Prob}(g_1, \dots, g_N) = \begin{cases} \frac{1}{D^{N-1}} & \text{if } g_1 \oplus \dots \oplus g_N = g, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

The measurement scenario involved in Mermin's original proof for the 3 qubit GHZ state involves the following four measurement contexts: XXX (the control), XYX , YXY and YYX (the variations). In the qubit case ($\mathbb{G} = \mathbb{Z}_2$), a measurement in the Pauli Y observable can equivalently be obtained by applying phase gate $P_{\pi/2}$ followed by a measurement in the X observable. Hence the measurement scenario above can equivalently be phrased in terms of the following Mermin measurements: $(0, 0, 0)$, $(0, \pi/2, \pi/2)$, $(\pi/2, 0, \pi/2)$ and $(\pi/2, \pi/2, 0)$.

Now we define Mermin measurement scenarios for our generalized setting. Take any consistent system of equations of \mathbb{Z} -modules valued in \mathbb{K} , i.e. one in the form

$$\left(\bigoplus_{m=1}^M n_m^l y_m = P_{a^l} \right)_{l=1}^L, \quad (8)$$

where y_1, \dots, y_M are variables, n_1^l, \dots, n_M^l are integers and $P_{a^l} \in \mathbb{K}$ (i.e. $a^l \in \mathbb{G}$). Assume that system (8) admits a solution $(y_m := P_{\beta_m})_{m=1}^M$ in the group \mathbb{P} of phase gates, and without loss of generality assume that a^1, \dots, a^L jointly span the entirety of \mathbb{G} [21]. By convention we set both $\beta_0 := 0_{\mathbb{G}}$ and $a^0 := 0_{\mathbb{G}}$. Let k be the exponent of \mathbb{K} ; choose $N > \sum_{m=1}^M n_m^l$ for all $l = 1, \dots, L$ and such that $N = 1 \pmod{k}$; define $n_0^l := N - \sum_{m=1}^M n_m^l$. For $i = 1, \dots, N$ consider phase gates $P_{\alpha_i^l} \in \mathbb{P}$ defined as follows:

- (i) define $\mu^l : \{1, \dots, N\} \rightarrow \{0, \dots, M\}$ by letting $\mu^l(i)$ be the least $\mu \geq 0$ such that $i \leq \sum_{m=0}^{\mu} n_m^l$;
- (ii) for $i = 1, \dots, N$, let $\alpha_i^l := \beta_{\mu^l(i)}$.

The *Mermin measurement scenario* associated with system (8) (together with a solution $(y_m := P_{\beta_m})_{m=1}^M$) is

given by taking the following $NL + 1$ Mermin measurements as the set of *measurement contexts*, which we shall henceforth denote by \mathcal{M} :

$$\left\{ \begin{array}{ll} (0, 0, \dots, 0, 0) & \text{the control} \\ (\alpha_1^1, \alpha_2^1, \dots, \alpha_{N-1}^1, \alpha_N^1) & \text{the 1st variation for } l = 1 \\ (\alpha_2^1, \alpha_3^1, \dots, \alpha_N^1, \alpha_1^1) & \text{the 2nd var. for } l = 1 \\ \vdots & \\ (\alpha_N^1, \alpha_1^1, \dots, \alpha_{N-2}^1, \alpha_{N-1}^1) & \text{the Nth var. for } l = 1 \\ \vdots & \\ (\alpha_1^L, \alpha_2^L, \dots, \alpha_{N-1}^L, \alpha_N^L) & \text{the 1st var. for } l = L \\ \vdots & \\ (\alpha_N^L, \alpha_1^L, \dots, \alpha_{N-2}^L, \alpha_{N-1}^L) & \text{the Nth var. for } l = L \end{array} \right. \quad (9)$$

From (7), the outcomes (g_1, \dots, g_N) for the control are uniformly distributed amongst those satisfying

$$g_1 \oplus \dots \oplus g_N = 0 \oplus \dots \oplus 0 = 0.$$

Similarly, the outcomes (g_1, \dots, g_N) for each of the variations are uniformly distributed amongst those satisfying

$$\begin{aligned} g_1 \oplus \dots \oplus g_N &= \alpha_1 \oplus \dots \oplus \alpha_N = \\ &= n_0^l \oplus n_1^l \beta_1 \oplus \dots \oplus n_M^l \beta_M = a^l. \end{aligned}$$

Contextuality of Mermin measurement scenarios is governed by the following result from [5].

Theorem 1. *A Mermin measurement scenario in the form of (9) is contextual if and only if system (8) admits no solution in the subgroup \mathbb{K} .*

In the original Mermin scenario, contextuality followed because the individual equation $2y = 1$ didn't have solutions in $\mathbb{Z}_2 \cong \mathbb{K}$. If contextual, a Mermin measurement scenario is an All-vs-Nothing argument [7]: therefore it is strongly contextual [22], and in particular maximally contextual [23]. Maximal contextuality is also known as maximal non-locality [24], or zero local fraction [25], and it means that a scenario lies on a face of the nonsignaling polytope having no local elements. It can further be shown that any consistent system admits a solution in \mathbb{P} , i.e. that all Mermin measurement scenarios are quantum realizable [7]; this leads us to propose their application to a QSS protocol.

Protocol design.—The following parameters are fixed beforehand and define the specific instance of the protocol to be executed.

- (a) A D -dimensional quantum system \mathcal{H} ; a group \mathbb{G} of order D and exponent k ; an orthonormal basis $|g\rangle_{g \in \mathbb{G}}$ of \mathcal{H} .
- (b) A system in the form of (8); a solution $(y_m := P_{\beta_m})_{m=1}^M$ in \mathbb{P} (let $\beta_0 := 0$); an appropriate N .

Suppose that the *dealer* (call her Alice) wishes to share S D -its (corresponding to $S \log_2 D$ bits) of a *secret* with $N' < N$ *players*. We wish to ensure that the secret can be decoded from the information Alice sends if and only if all players agree to cooperate (by which we mean that they all reveal their *secret keys* to some party in possession of the information). We assume the following *security conditions* to hold.

- (i) Alice and the players share an authenticated classical channel, ensuring integrity and availability of all classical communications involved in the protocol. Confidentiality is not guaranteed.
- (ii) Alice and the players are in possession of N secure independent classical sources of randomness which she can trust, to generate *measurement choices* valued in $\{0, 1, \dots, M\}$.
 - (iia) Alice is in possession of a secure classical source of randomness which she can trust, independent from all other. This will be used to decide which rounds will be *secret rounds*, with probability $\sigma > 0$, and which rounds will be *test rounds*, with probability $\tau = 1 - \sigma > 0$.
 - (iib) Alice is in possession of a secure classical source of randomness which she can trust, independent from all other. This will be used to decide which rounds will be *secret rounds*, with probability $\sigma > 0$, and which rounds will be *test rounds*, with probability $\tau = 1 - \sigma > 0$.
 - (iic) For the whole execution of step 2 of the protocol below, no signaling is possible between distinct players or devices. This can be achieved by ensuring the devices are operated in conditions controlled by Alice (trusted laboratories, synchronized timestamp servers and bit commitment, etc).
 - (iiv) The issue of players communicating different measurement choices from the ones they actually executed is already covered in [1], so we will assume that in step 3 Alice is communicated the measurement choices faithfully. This can be achieved by entrusting the laboratory setup with the communication of the random measurement choices to both the measuring device and Alice, with restricted player control over it.

The device-independent quantum secret sharing (DIQSS) protocol proceeds as follows for each round, until the entire secret has been transmitted. The setup is illustrated in Fig. (2).

1. Alice and the players share N subsystems of a state ψ : each player has an individual subsystem and Alice keeps the remaining $N - N'$ subsystems. In a noiseless, trusted implementation, ψ is the pure quantum state $|\text{GHZ}_{0\mathbb{G}}^N\rangle$. For the purposes of our device-independent security proof, ψ can be any state of some possibly postquantum theory, and an eavesdropper (call her Eve) may keep additional subsystems to herself.

2. Alice and the players each sample their classical source of randomness and obtain measurement choices m_1, \dots, m_N which are given as input to some *measuring devices* B_1, \dots, B_N and result in outputs $g_1, \dots, g_{N'}$ for the players (the *secret keys*) and $g_{N'+1}, \dots, g_N$ for Alice. In a noiseless, trusted implementation, B_j with input m_j applies the phase gate $P_{\beta_{m_j}}$ to the subsystem j and then measures it in the generalized X observable.
3. The measurement choices are communicated to Alice. She checks that (m_1, \dots, m_N) defines a valid measurement context, and otherwise skips to the next round. Refer to the Appendix for further discussion of the failure rate of this step, together with possible mitigation measures.
4. Alice samples her source of randomness to decide whether the round will be a test or a secret round.
 - 4a. If the round is a test round, Alice requests all players to communicate their secret keys $g_1, \dots, g_{N'}$. If the measurement context is the control, she checks that $g_1 \oplus \dots \oplus g_N = 0$. If the measurement context is a variation for some $l = 1, \dots, L$, she checks that $g_1 \oplus \dots \oplus g_N = a^l$. If the check fails, she records the failure and skips to the next round.
 - 4b. If the round is a secret round, Alice computes $g_{dealer} := \bigoplus_{j=N'+1}^N g_j$ and communicates the ciphertext $c := s \oplus g_{dealer}$ to the players, where the plaintext s is the next D -it of the secret.
5. Anyone in possession of the ciphertext c and all secret keys $g_1, \dots, g_{N'}$ for a certain round can obtain the plaintext s by computing $s = (c \oplus g_1 \oplus \dots \oplus g_{N'}) \ominus a^l$, where $l = 0$ if the measurement context was the control, and ranges over $1, \dots, L$ for the variations.

Confidentiality and security.—A trusted implementation of a quantum secret sharing protocol typically comes with two assurances: it must guarantee that ignorance about any one secret key denies knowledge about the plaintext, and it must guarantee that successful eavesdropping has low enough probability. For the original HBB QSS protocol, the first guarantee follows abstractly from strong complementarity of the Pauli X and Z observables [3], and generalizes to our protocol. The second guarantee originally followed from mutual unbiased of the X and Y observables, and a more careful analysis is needed here. It can be shown (see Appendix) that the optimal eavesdropping probability $P(\text{eavesdrop})$ can be obtained from the Rényi $\alpha = 2$ entropy of $\rho := \sum_{m=0}^M P(m) |\beta_m\rangle \langle \beta_m|$ as follows:

$$P(\text{eavesdrop}) = 2^{-H_2(\rho)},$$

where $|\beta_m\rangle$ is the phase state for phase gate P_{β_m} .

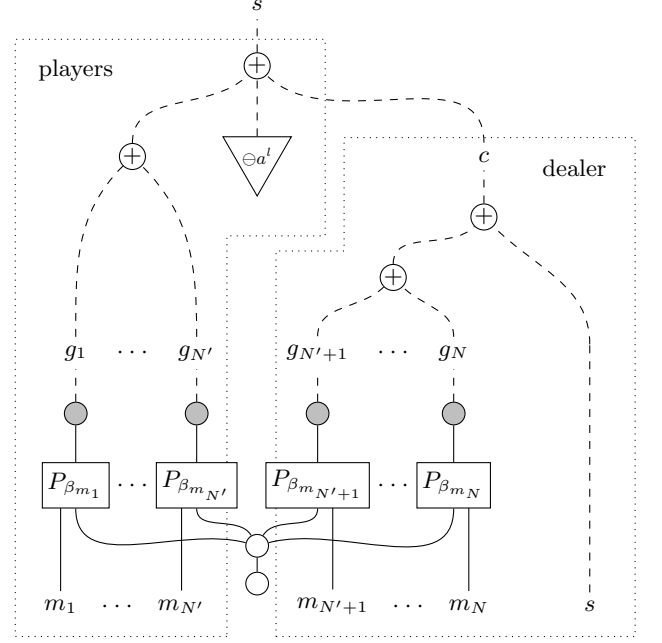


FIG. 2. Graphical presentation of a noiseless, trusted implementation for the DIQSS protocol. The shared generalized GHZ state is denoted in white; measurements in the generalized X observable are denoted in gray; phase gates are denoted by boxes, and controlled by the classical variables m_1, \dots, m_N . The \oplus nodes denote classical addition in \mathbb{G} , and $\ominus a^l$ is the inverse of a^l in \mathbb{G} .

These results don't require contextuality, and are only reliable for a trusted implementation. Assuming that the conditions guaranteeing contextuality in Theorem (1) hold, we now show device-independent security for the DIQSS protocol. We start by conditioning the state available to Alice and the players on any classical information e that Eve can obtain from side-channels (i.e. by measuring some subsystem of state ψ she may have kept to herself), as well as on any classical information i that could have been used in the preparation of state ψ . We obtain some ensemble $\sum_{ie} q^{ie} \psi^{ie}$. Each state ψ^{ie} , together with the black boxes and the nonsignaling condition, gives rise to an empirical model $(e_C^{ie})_{C \in \mathcal{M}}$, where \mathcal{M} contains those combinations (m_1, \dots, m_N) of measurement choices, the *measurement contexts*, which don't result in the protocol being aborted by Alice in step 3.

We will say that an empirical model passes the *Mermmin contextuality test* (MCT) in a context $C \in \mathcal{M}$ if it would pass a test round with certainty when the measurement context checked in step 3 is C . We will say that an empirical model passes the *context-independent MCT* if it passes the MCT with certainty for all contexts. Each empirical model e^{ie} can be decomposed as $e^{ie} = \delta^{ie} \xi^{ie} + (1 - \delta^{ie}) \zeta^{ie}$, for some unique empirical models ζ^{ie} and ξ^{ie} and some $\delta^{ie} \in [0, 1]$, where ζ^{ie} passes the context-independent MCT and ξ^{ie} fails test rounds with certainty in at least one measurement context.

The empirical model ζ^{ie} passing the context-independent MCT can in turn be written as a convex combination $\zeta^{ie} = \sum_t p_t^{ie} \zeta[t]$ of certain empirical models $\zeta[t]$, defined in the Appendix and independent of i and e . The empirical models $\zeta[t]$ have the property that, whenever $\{1, \dots, N\} = I \sqcup J$ is a partition with $2 \leq |I| \leq N$, the conditional distribution $\zeta[t](g_i)_{i \in I} | (g_j)_{j \in J}$ is uniform over some set of size $D^{|I|-1}$. (The relevance of the $|I| \geq 2$ requirement stems from the fact that at least two measurement outcomes are not revealed in an attack: at least one outcome in the hands of Alice, and one outcome in the hands of at least one non-cooperating player.)

Because of the uniformity of the $\zeta[t]$ outcomes, as long as g_i and $g_{i'}$ remain unknown for two distinct i, i' we have that, with probability $(1 - \delta^{ie})$, the state ψ^{ie} carries at least one D -it of randomness with respect to: (i) any past information i ; (ii) any information e that Eve may have obtained from the state ψ^{ie} ; (iii) knowledge of all other g_k for $k \neq i, i'$. Now assume that one player does not cooperate and that the plaintext carries one D -it of randomness: then with probability $(1 - \delta^{ie})$ Eve cannot use ψ^{ie} to carry any information about past plaintexts (via i) or the current plaintext (via e). Therefore with probability $(1 - \delta) := \sum_{ie} q^{ie} (1 - \delta^{ie})$ we have that Eve is physically forbidden from performing any memory attack or eavesdropping, simply by requiring context-independent success of the MCT. On the other hand, with probability $\delta := \sum_{ie} q^{ie} \delta^{ie}$ we have that Eve — whom we have assumed to possess potentially postquantum capabilities — can use the nonsignaling empirical models ξ^{ie} to do whatever she wants with i and the current secret, at the price of ξ^{ie} failing the MCT in at least one context.

After all S D -its of the secret have been transmitted, Alice checks the observed failure rate ϵ during test rounds. If this is lower than some $\epsilon_{max} \ll 1$ which she believes to be reasonable for a noisy trusted implementation, the protocol has succeeded and the secret can be used; otherwise the protocol has failed. With probability at least p_{min} , the smallest probability amongst all measurement contexts, Alice would have discovered one of the ξ^{ie} empirical models, so the expected contribution of Eve's machinations to the total failure rate ϵ is necessarily bounded above by $p_{min} \delta$. Therefore the amount of information E_{total} (in D -its) learned by Eve on the secret is bounded as follows with high probability:

$$E_{total} \leq \frac{1}{p_{min}} \left[\epsilon + O\left(\sqrt{\epsilon/T}\right) \right] S.$$

where the big O notation hides a small constant, and T is the number of test rounds that have been executed (expected $T \approx \frac{T}{\epsilon} S$). This bound is minimized by choosing a distribution as uniform as possible over measurement contexts, approaching the optimal $p_{min} = \frac{1}{|\mathcal{M}|}$.

Because of the randomness required for the plaintexts and the possible leaking of information — associated with a non-zero acceptable failure rate during test rounds, or

with failure of the protocol altogether — the secret being shared should be a disposable ephemeral key. Once the protocol has succeeded, Alice can encrypt and broadcast the actual secret, which cooperating players can decrypt using the key they now share.

Discussion.—We have presented a new device independent quantum secret sharing (DIQSS) protocol, deriving its correctness from a group theoretic understanding of generalized GHZ states. Recent developments in the characterization of Mermin contextuality scenarios have allowed us to prove security of the protocol against nonsignaling postquantum attackers. We showed that a secret can be securely shared with a maximum fraction of leaked information peaked around $\epsilon_{max} |\mathcal{M}|$, a device-independent bound which only relies on:

- (i) the maximum failure rate ϵ_{max} for test rounds that is acceptable to the dealer, and
- (ii) the number $|\mathcal{M}|$ of measurement contexts required to prove Mermin contextuality.

The bound is obtained by exploiting a particularly strong form of contextuality shown by Mermin measurement scenarios. We decomposed the empirical models associated with specific implementations of the protocol into the component passing test rounds with certainty and the component failing test rounds with certainty for some choice of measurement context. We showed that, because of maximal contextuality, the component passing test rounds with certainty is uniformly random with respect to any side-channel information and past information, and hence cannot be used to eavesdrop or perform memory attacks. The other component can be used to perform attacks, but at the cost of a quantifiable increase in failure rate for test rounds. In a nutshell, the maximal contextuality of Mermin measurement scenarios means that an attacker can exploit the acceptable failure rate for contextuality tests (quantified by ϵ_{max}) and the difficulty of proving contextuality (quantified by $|\mathcal{M}|$), but nothing more than that.

We conclude with a handful of open questions. Firstly, the failure rate of step 3 is quite high: a rough analysis and some mitigation proposals can be found in the Appendix. Secondly, the $\epsilon_{max} |\mathcal{M}|$ bound on the leaked fraction of the secret is quite lax: not all the measurement contexts are needed in Mermin measurement scenarios to prove contextuality (in the general case), and a more thorough analysis may result in a lower bound. Finally, it would be interesting to have an extension of proof of security against postquantum attackers which includes standard privacy amplification techniques. We hope that future work will address these points.

The author would like to thank Bob Coecke, Jonathan Barrett, Matty Hoban and Nathan Walk for their comments and suggestions. Funding from EPSRC and Trinity College is gratefully acknowledged.

Appendix A: eavesdropping probability

* stefano.gogioso@cs.ox.ac.uk

- [1] M. Hillery, V. Bužek, and A. Berthiaume, *Physical Review A* **59**, 1829 (1999).
- [2] L. Xiao, G. L. Long, F.-G. Deng, and J.-W. Pan, *Physical Review A* **69** (2004), 10.1103/PhysRevA.69.052307.
- [3] V. N. Zamdzhiev, *An abstract approach towards quantum secret sharing*, Master's thesis, University of Oxford (2012).
- [4] N. D. Mermin, *Am. J. Phys* **58**, 731 (1990).
- [5] S. Gogioso and W. Zeng, (2015), arXiv:1506.02675.
- [6] B. Coecke, R. Duncan, A. Kissinger, and Q. Wang (IEEE Computer Society, 2012).
- [7] S. Gogioso, "Operational mermin non-locality and all- vs-nothing arguments," (2015).
- [8] J. Barrett, L. Hardy, and A. Kent, *Physical Review Letters* **95**, 010503 (2005).
- [9] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani, *Physical Review Letters* **98**, 230501 (2007).
- [10] K. Horodecki, M. Horodecki, P. Horodecki, R. Horodecki, M. Pawłowski, and M. Bourennane, (2010), arXiv:1006.0468.
- [11] U. Vazirani and T. Vidick, *Physical review letters* **113**, 140501 (2014).
- [12] B. Coecke and É. O. Paquette, in *New Structures for Physics* (Springer, 2011) pp. 173–286.
- [13] P. Selinger, in *New structures for physics* (Springer, 2011) pp. 289–355.
- [14] S. Abramsky and B. Coecke, (2008), 10.1016/B978-0-444-52869-8.50010-4.
- [15] S. Abramsky and B. Coecke, in *Logic in Computer Science, 2004.* (IEEE, 2004) pp. 415–425.
- [16] A. Joyal and R. Street, *Advances in Mathematics* **88**, 55 (1991).
- [17] A. Joyal and R. Street, *Advances in Mathematics* **102**, 20 (1993).
- [18] B. Coecke and S. Perdrix, in *Computer Science Logic* (Springer, 2010) pp. 230–244.
- [19] P. Selinger, *ENTCS* **170**, 139 (2007).
- [20] A. Kissinger, *Pictures of processes: automated graph rewriting for monoidal categories and applications to quantum computing*, Ph.D. thesis (2012), arXiv:1203.0202.
- [21] If they don't, then the argument can be formulated in the smaller abelian group defined by their span, and the security of the protocol is given by the order D' of this smaller group, instead of the order D of \mathbb{G} .
- [22] S. Abramsky, R. S. Barbosa, K. Kishida, R. Lal, and S. Mansfield, (2015), arXiv:1502.03097.
- [23] S. Abramsky and A. Brandenburger, *New Journal of Physics* **13**, 113036 (2011).
- [24] L. Aolita, R. Gallego, A. Acín, A. Chiuri, G. Vallone, P. Mataloni, and A. Cabello, *Physical Review A* **85**, 032107 (2012).
- [25] A. C. Elitzur, S. Popescu, and D. Rohrlich, *Physics Letters A* **162**, 25 (1992).
- [26] Recall that, without loss of generality, we have assumed $\mathbb{H} = \mathbb{G}$, which yields the result quoted in the main body.
- [27] H.-K. Lo, H. F. Chau, and M. Ardehali, *Journal of Cryptology* **18**, 133 (2005).

Suppose that Eve measures one of the subsystems in the X^γ observable, by which we mean that she applies some phase P_γ and then performs a (non-demolition) measurement in the generalized X observable, obtaining outcome $h \in \mathbb{G}$. We want to compute the probability $P(\text{eavesdrop})$ of Eve guessing the output. This apparently simple scenario, in which success is defined by Eve learning one secret key, is extremely relevant in this setting: one secret key is all which a group of $N'-1$ colluding players needs to circumvent the confidentiality guarantee of the protocol. Suppose that, after Eve has secretly measured the state obtaining $h \in \mathbb{G}$, the player measures in the X^{β_m} observable, with some probability $P(m)$ for each $m = 0, 1, \dots, M$, and obtains outcome $g \in \mathbb{G}$. The probability $P(\text{eavesdrop}) := P(g=h)$ that Eve successfully learned the secret key is:

$$P(\text{eavesdrop}) = \sum_{m=1}^M P(m) \left| \frac{1}{D} \sum_{\chi \in G^\wedge} e^{i(\beta_m(\chi) - \gamma(\chi))} \right|^2$$

The term in square absolute value is obtained by observing that outcome h for the observable X^γ is obtained by measuring $P_{-\gamma}|h\rangle$, and then expanding $\langle g|P_{\beta_m}P_{-\gamma}|g\rangle$ as

$$\frac{1}{D^2} \sum_{\chi, \chi' \in G^\wedge} e^{i(\beta_m(\chi) - \gamma(\chi'))} \langle g|\chi\rangle \langle \chi'|g\rangle.$$

Define phase state $|\alpha\rangle$ associated to phase gate P_α to be

$$|\alpha\rangle := \frac{1}{\sqrt{D}} \sum_{\chi \in G^\wedge} e^{i\alpha(\chi)} \frac{1}{\sqrt{D}} |\chi\rangle$$

and rewrite the eavesdropping probability as

$$P(\text{eavesdrop}) = \sum_{m=1}^M P(m) |\langle \beta_m | \gamma \rangle|^2 = \langle \gamma | \rho | \gamma \rangle, \quad (10)$$

where ρ is the following mixed state:

$$\rho := \sum_{m=0}^M P(m) |\beta_m\rangle \langle \beta_m|.$$

The optimal eavesdropping probability $P(\text{eavesdrop})$ can then be related to the Rényi $\alpha = 2$ entropy (or collision entropy) of the mixed state ρ as

$$P(\text{eavesdrop}) = 2^{-H_2(\rho)}, \quad (11)$$

and the optimal eavesdropping strategy, where ρ is diagonalized as $\sum_i q_i |i\rangle \langle i|$, is to choose $|\gamma\rangle = \sum_i q_i |i\rangle$. In the case where phases are mutually unbiased and $P(m)$ is the uniform distribution, the probability of eavesdropping is $P(\text{eavesdrop}) = \frac{1}{M+1} + \frac{M}{D(M+1)}$, which approximates uniformly random guessing for high dimension $D \gg M$.

Appendix B: Mermin contextuality tests

We show that the empirical models satisfying the context-independent Mermin contextuality test (MCT) are exactly those that can be obtained as convex combinations of the empirical models $\zeta[t]$ defined below, where $t \equiv (t_j)_{j=1}^N \in \mathbb{G}^N$. Let \mathbb{H} be the subgroup of \mathbb{G} spanned by a^1, \dots, a^L [26], and let $\mathbb{F}_g < \mathbb{G}^N$ be the subgroup

$$\mathbb{F}_g := \{(g_1, \dots, g_N) \in \mathbb{G}^N \mid g_1 \oplus \dots \oplus g_N = g\}.$$

We will show that the distribution $\zeta[t]_C(g)$, for $g \equiv (g_1, \dots, g_N) \in \mathbb{G}^N$, is necessarily given by:

$$\zeta[t]_C(g) = \begin{cases} \frac{1}{|\mathbb{H}|^{N-1}} & \text{if } C \text{ is the control,} \\ & \text{and } g \in (t \oplus \mathbb{H}^N) \cap \mathbb{F}_{0_{\mathbb{G}}} \\ \frac{1}{|\mathbb{H}|^{N-1}} & \text{if } C \text{ is a variation (for some } l), \\ & \text{and } g \in (t \oplus \mathbb{H}^N) \cap \mathbb{F}_{a^l} \\ 0 & \text{otherwise} \end{cases}$$

Take any empirical model ζ which minimally satisfies the context-independent MCT, i.e. such that it cannot be obtained as a convex combination of models satisfying it. The empirical model ζ is maximally contextual, as a consequence of Theorem (1). Now consider the possibilistic empirical model $b \cdot \zeta$, where $b : (\mathbb{R}^+, +, 0, \times, 1) \rightarrow (\mathbb{B}, \vee, 0, \wedge, 1)$ is the semiring homomorphism from the positive reals to the booleans sending 0 to 0 and any $x > 0$ to 1; for each measurement context $C \in \mathcal{M}$, $(b \cdot \zeta)_C$ is the indicator function for the support of probability distribution ζ_C . Because ζ is maximally contextual, it is in particular possibilistically contextual [23], i.e. $b \cdot \zeta$ is contextual over the booleans. (Possibilistic contextuality is strictly stronger than contextuality, and strictly weaker than maximal contextuality.) In fact, $b \cdot \zeta$ necessarily satisfies the context-independent MCT over the booleans, from which we deduce the following.

- (i) For some $t = (t_1, \dots, t_N) \in \mathbb{F}_{0_{\mathbb{G}}}$, we have $(b \cdot \zeta)_{\text{control}}(t) = 1$. Relabel ζ as $\zeta[t]$.
- (ii) By nonsignaling, if C is any variation (for some $l = 1, \dots, L$), then $(b \cdot \zeta[t])_C(t_{j,a^l}) = 1$ for all $j = 1, \dots, L$, where t_{j,a^l} is defined by:

$$(t_{j,a^l})_i := \begin{cases} t_i & \text{if } i \neq j \\ t_i \oplus a_l & \text{if } i = j \end{cases}$$

- (iii) By nonsignaling, $(b \cdot \zeta[t])_{\text{control}}(t_{j,j',a^l}) = 1$ for all distinct j, j' , where t_{j,j',a^l} is defined by:

$$(t_{j,j',a^l})_i := \begin{cases} t_i & \text{if } i \neq j, j' \\ t_i \oplus a_l & \text{if } i = j \\ t_i \ominus a_l & \text{if } i = j' \end{cases}$$

By repeating steps (ii)-(iii) above, we get that $(b \cdot \zeta[t])_{\text{control}}$ must be 1 over all elements of $(t \oplus \mathbb{H}^N) \cap \mathbb{F}_{0_{\mathbb{G}}}$, and that any variation C (for some value of l) must be 1 over all elements of $(t \oplus \mathbb{H}^N) \cap \mathbb{F}_{a^l}$. With only t as a starting element, this is the most that can be obtained by enforcing context-independent MCT and nonsignaling, and $(b \cdot \zeta[t])$ is minimal satisfying the context-independent MCT over the booleans. The last thing we have to show is that $\zeta[t]$, as a probabilistic empirical model, is uniformly distributed over its support. By nonsignaling we must have $\zeta[t]_{\text{control}}(t) = \zeta[t]_C(t_{j,a^l})$ for all variations C and all j , and similarly $\zeta[t]_C(t_{j,a^l}) = \zeta[t]_{\text{control}}(t_{j,j',a^l})$ for all $j' \neq j$. Because the entire support of $\zeta[t]$ was generated via steps (ii)-(iii) above, these equalities are sufficient to conclude our proof.

Appendix C: failure rate for step 3

The simplified implementation presented here suffers from a low success rate for step 3: this problem is already known for the HBB QSS protocol [2, 27], but the mitigation proposed in [2, 27], biasing the probability of different measurement contexts, impacts the security of the DIQSS protocol. Because of the way Mermin contextuality arguments are set up, there are $LN + 1$ measurement contexts involved in a scenario, out of $(M + 1)^N$ possible measurement choices. However, there is no real need to restrict the variations to the cyclic permutations shown by (9), and any permutation $(\beta_{m_1}, \dots, \beta_{m_N})$ of a measurement context will have the same effect as the context itself. For an individual equation in the form $ky = P_a$ (here $M = 1$ and $L = 1$), this brings the success rate for step 3 from $\frac{N+1}{2N}$ to at least $\frac{1}{kD}$. Further improvement is possible. For example, let $\mathbb{G} = \mathbb{Z}_{p^r}$, set $k := p$ and consider a convex combination (with equal coefficients) of the p^{r-1} Mermin measurement scenarios for equations $(nky = n)_n$, for all $n \in \{1, \dots, p^{r-1} - 1\}$ not divisible by p . This increases the success probability for step 3 to $\frac{p-1}{pk}$, which is as high as $\frac{1}{4}$ for the $p = 2$ case.

Naively, one could think that the increase in acceptable measurement choices, and consequential decrease of p_{\min} , would inevitably lead to an increased amount of leaked information for fixed ϵ_{\max} . However, the security analysis above relies on the unspoken assumption that all measurement contexts are necessary to prove contextuality, an assumption that certainly fails for the extended scenarios suggested here. In fact, we don't expect the amount of leaked information to change at all. Also, it is likely that not all measurement contexts will be needed, in a general Mermin measurement scenario, to prove contextuality: confirmation of this would reduce the upper bound on the fraction of leaked information from $|\mathcal{M}|_{\epsilon_{\max}}$ to some smaller multiple of ϵ_{\max} .

Appendix D: perfect noncontextual attack

We provide a perfect attack in case the conditions guaranteeing contextuality in Theorem (1) were to fail, i.e. in case there was a solution $(y_m := P_{b_m})_{m=1}^M$ in \mathbb{K} to the chosen system of equations. In this case, Eve independently samples many times a $2N$ -partite classical mixture $(G_1, \dots, G_N, G_{N+1}, \dots, G_{2N})$ of \mathbb{G} values with the following properties:

- probability $P(G_1 = h_1, \dots, G_N = h_N)$ is $\frac{1}{D^{N-1}}$ if $h_1 \oplus \dots \oplus h_N = 0_{\mathbb{G}}$ and vanishes otherwise;
- conditional probability $P(G_j | G_{N+j})$ is 1 for all j .

She then encodes in each black box B_j the sequence of values sampled from G_j , for $j = 1, \dots, N$, and keeps the corresponding sequences for G_{N+1}, \dots, G_{2N} to herself. Each black box B_j is designed to take an input m_j , get the next stored sample h_j for G_j and return $g_j := h_j \oplus b_{m_j}$ as output. From the measurements $m_1, \dots, m_{N'}$ that the players send to Alice, Eve can reconstruct all secret keys $g_1, \dots, g_{N'}$ by looking up the relevant samples $h_{N+1} = h_1, \dots, h_{2N} = h_N$ of G_{N+1}, \dots, G_{2N} . This attack is undetectable, because it is designed to have the exact same statistics of the trusted protocol for all measurement contexts. As long as Eve can keep track of how many measurements have been made, she can learn all the plaintexts with Alice being none the wiser.